Punishing People or Learning from Failure?

The choice is ours

Sidney Dekker

Associate Professor

Centre for Human Factors in Aviation, IKP

Linköping Institute of Technology

SE - 581 83 Linköping

Sweden

tel. + 46 13 281646

fax. + 46 13 282579

sidde@ikp.liu.se

Disinheriting Fitts and Jones '47

Abstract

In this paper I describe how Fitts and Jones laid the foundation for aviation human factors by trying to understand why human errors made sense given the circumstances surrounding people at the time. Fitts and Jones remind us that human error is not the cause of failure, but a symptom of failure, and that "human error"—by any other name or by any other human—should be the starting point of our investigations, not the conclusion. Although most in aviation human factors embrace this view in principle, practice often leads us to the old view of human error which sees human error as the chief threat to system safety. I discuss two practices by which we quickly regress into the old view and disinherit Fitts and Jones: (1) the punishment of individuals, and (2) error classification systems. In contrast, real progress on safety can be made by understanding how people create safety, and by understanding how the creation of safety can break down in resource-limited systems that pursue multiple competing goals. I argue that we should de-emphasize the search for causes of failure and concentrate instead on mechanisms by which failure succeeds, by which the creation of safety breaks down.

Keywords: human error, mechanisms of failure, safety culture, human factors, classification, creation of safety

Running head: Disinheriting Fitts and Jones '47

Introduction

The groundwork for human factors in aviation lies in a couple of studies done by Paul Fitts and his colleague Jones right after World War II. Fitts and Jones (1947) found how features of World War II airplane cockpits systematically influenced the way in which pilots made errors. For example, pilots confused the flap and gear handles because these typically looked and felt the same and were co-located. Or they mixed up the locations of throttle, mixture and propeller controls because these kept changing across different cockpits. Human error was the starting point for Fitts' and Jones' studies—not the conclusion. The label "pilot error" was deemed unsatisfactory, and used as a pointer to hunt for deeper, more systemic conditions that led to consistent trouble. The idea these studies convey to us is that mistakes actually make sense once we understand features of the engineered world that surrounds people. Human errors are systematically connected to features of people's tools and tasks. The insight, at the time as it is now, was profound: the world is not unchangeable; systems are not static, not simply given. We can re-tool, re-build, re-design, and thus influence the way in which people perform. This, indeed, is the historical imperative of human factors—understanding why people do what they do so we can tweak, change the world in which they work and shape their assessments and actions accordingly.

Years later, aerospace human factors extended the Fitts and Jones work. Increasingly, we realized how trade-offs by people at the sharp end are influenced by what happens at the blunt end of their operating worlds; their organizations (Maurino *et al.*, 1995). Organizations make resources available for people to use in local workplaces (tools, training, teammates) but put constraints on what goes on there at the same time (time pressures, economic considerations), which in turn influences the way in which people decide and act in context (Woods *et al.*, 1994; Reason, 1997). Again, what people do makes sense on the basis of the circumstances surrounding them, but now circumstances that reach far beyond their immediate engineered interfaces. This realization has put the Fitts and Jones premise to work in organizational contexts, for example changing workplace conditions or reducing working hours or de-emphasizing production to encourage safer trade-offs on the line (e.g. the "no fault go-around policy" held by many airlines today, where no (nasty) questions will be asked if a pilot breaks off his attempt to land). Human error is still systematically connected to features of people's tools and tasks, and, as acknowledged more recently, their operational and organizational environment.

Two views of human error

These realizations of aviation human factors pit one view of human error against another. In fact, these are two views of human error that are almost totally irreconcilable. If you believe one or pursue countermeasures on its basis, you truly are not able to embrace the tenets and putative investments in safety of the other. The two ways of looking at human error are that we can see human error as a cause of failure, or we can see human error as a symptom of failure (Woods *et al.*, 1994). The two views have recently been characterized as *the old view* of human error versus *the new view* (Cook, Render & Woods, 2000; AMA, 1998; Reason, 2000) and painted as fundamentally irreconcilable perspectives on the human contribution to system success and failure. In the old view of human error:

- Human error is the cause of many accidents.
- The system in which people work is basically safe; success is intrinsic. The chief threat to safety comes from the inherent unreliability of people.
- Progress on safety can be made by protecting the system from unreliable humans through selection, proceduralization, automation, training and discipline.

This old view was the one that Fitts and Jones remind us to be skeptical of. Instead, implicit in their work was the new view of human error:

- Human error is a symptom of trouble deeper inside the system.
- Safety is not inherent in systems. The systems themselves are contradictions between multiple goals that people must pursue simultaneously. People have to create safety.
- Human error is systematically connected to features of peoples tools, tasks and operating environment. Progress on safety comes from understanding and influencing these connections.

Perhaps everyone in aviation human factors wants to pursue the new view. And most people and organizations certainly posture as if that is exactly what they do. Indeed, it is not difficult to find proponents of the new view—in principle—in aerospace human factors. For example:

"...simply writing off aviation accidents merely to pilot error is an overly simplistic, if not naive, approach.... After all, it is well established that accidents cannot be attributed to a single cause, or in most instances, even a single individual. In fact, even the identification of a 'primary' cause is fraught with problems. Instead, aviation accidents are the result of a number of causes..." (Shappell & Wiegmann, 2001, p. 60).

In practice, however, attempts to pursue the causes of system failure according to the new view can become retreads of the old view of human error. In practice, getting away from the tendency to judge instead of explain turns out to be difficult; avoiding the fundamental attribution error remains very hard; we tend to blame the man-in-the-loop. This is not because we aim to blame—in fact, we probably intend the opposite. But roads that lead to the old view in aviation human factors are paved with intentions to follow the new view. In practice, we all too often choose to disinherit Fitts and Jones '47, frequently without even knowing it. In this paper, I try to shed some light on how this happens, by looking at the pursuit of individual culprits in the wake of failure, and at error classification systems. I then move on to the new view of human error, extending it with the idea that we should de-emphasize the search for causes and instead concentrate on understanding and describing the mechanisms by which failure succeeds.

The Bad Apple Theory I: Punish the culprits

Progress on safety in the old view of human error relies on selection, training and discipline—weeding and tweaking the nature of human attributes in complex systems that themselves are basically safe and immutable. For example, Kern (1999) characterizes "rogue pilots" as extremely unreliable elements, which the system, itself safe, needs to identify and contain or exile:

"Rogue pilots are a silent menace, undermining aviation and threatening lives and property every day.... Rogues are a unique brand of undisciplined pilots who place their own egos above all else—endangering themselves, other pilots and their passengers, and everyone over whom they fly. They are found in the cockpits of major airliners, military jets and in general aviation...just one poor decision or temptation away from fiery disaster." (back cover)

The system, in other words, contains bad apples. In order to achieve safety, it needs to get rid of them, limit their contribution to death and destruction by discipline, training or taking them to court (e.g. Wilkinson, 1994). In a recent comment, *Aviation Week and Space Technology* (North, 2000) discusses Valujet 592 which crashed after take-off from Miami airport because oxygen generators in its forward cargo hold had caught fire. The generators had been loaded onto the airplane without shipping caps in place, by employees of a maintenance contractor, who were subsequently prosecuted. The editor:

"...strongly believed the failure of SabreTech employees to put caps on oxygen generators constituted willful negligence that led to the killing of 110 passengers and crew. Prosecutors were right to bring charges. There has to be some fear that not doing one's job correctly could lead to prosecution." (p. 66)

Fear as investment in safety? This is a bizarre notion. If we want to know how to learn from failure, the balance of scientific opinion is quite clear: fear doesn't work. In fact, it corrupts opportunities to learn. Instilling fear does the opposite of what a system concerned with safety really needs: learn from failure by learning about it before it happens. This is what safety cultures are all about: cultures that allow the boss to hear bad news. Fear stifles the flow of safety-related information—the prime ingredient of a safety culture (Reason, 1997). People will think twice about going to the boss with bad news if the fear of punishment is hanging over their heads. Many people believe that we can punish and learn at the same time. This is a complete illusion. The two are mutually exclusive. Punishing is about keeping our beliefs in a basically safe system intact. Learning is about changing these beliefs, and changing the system. Punishing is about seeing the culprits as unique parts of the failure. Learning is about seeing the failure as a part of the system. Punishing is about stifling the flow of safety-related information. Learning is about increasing that flow. Punishing is about closure, about moving beyond the terrible event. Learning is about continuity, about the continuous improvement that comes from firmly integrating the terrible event in what the system knows about itself. Punishing is about not getting caught the next time. Learning is about countermeasures that remove error-producing conditions so there won't be a next time.

The construction of cause

Framing the cause of the Valujet disaster as the decision by maintenance employees to place unexpended oxygen generators onboard without shipping caps in place immediately implies a *wrong* decision, a *missed* opportunity to prevent disaster, a *disregard* of safety rules and practices. Framing of the cause as a decision leads to the identification of responsibility of people who made that decision which in turns leads to the legal pursuit of them as culprits. The Bad Apple Theory reigns supreme. It also implies that cause can be found, neatly and objectively, in the rubble. The opposite is true. We don't find causes. We construct cause. "Human error", if there were such a thing, is not a question of individual single-point failures to notice or process—not in this story and probably not in any story of breakdowns in flight safety. Practice that goes sour spreads out over time and in space, touching all the areas that usually make practitioners successful. The "errors" are not surprising brain slips that we can beat out of people by dragging them before a jury. Instead, errors are series of actions and assessments that are systematically connected to people's tools and tasks and environment; actions and assessments that often make complete sense when viewed from inside their situation. Were one to trace "the cause" of failure, the causal network would fan out immediately, like cracks in a window, with only the investigator determining when to stop looking because the evidence will not do it for him or her. There is no single cause. Neither for success, nor for failure.

The SabreTech maintenance employees inhabited a world of boss-men and sudden firings. It was a world of language difficulties—not just because many were Spanish speakers in an environment of English engineering language, as described by Langewiesche (1998, p. 228):

"Here is what really happened. Nearly 600 people logged work time against the three Valujet airplanes in SabreTech's Miami hangar; of them 72 workers logged 910 hours across several weeks against the job of replacing the "expired" oxygen generators—those at the end of their approved lives. According to the supplied Valujet work card 0069, the second step of the seven-step process was: 'If the generator has not been expended install shipping cap on the firing pin.' This required a gang of hard-pressed mechanics to draw a distinction between canisters that were 'expired', meaning the ones they were removing, and canisters that were not 'expended', meaning the same ones, loaded and ready to fire, on which they were now expected to put nonexistent caps. Also involved were canisters which were expired and expended, and others which were not expired but were expended. And then, of course, there was the simpler thing—a set of new replacement canisters, which were both unexpended and unexpired."

And, oh by the way, as you may already have picked up: there *were* no shipping caps to be found in Miami. How can we prosecute people for not installing something we do not provide them with? The pursuit of culprits disinherits the legacy of Fitts and Jones. One has to side with Hawkins (1987, p. 127) who argues that exhortation (via punishment, discipline or whatever measure) "is unlikely to have any long-term effect unless the exhortation is accompanied by other measures... A more profound inquiry into the nature of the forces which drive the activities of people is necessary in order to learn whether they can be manipulated and if so, how". Indeed, this was Fitts's and Jones's insight all along. If researchers could understand and modify the situation in which humans were required to perform, they could understand and modify the performance that went on inside of it. Central to this idea is the local rationality principle (Simon, 1969; Woods *et al.*, 1994). People do reasonable, or locally rational things given their tools, their multiple goals and pressures, their knowledge and their limited resources. Human error is a symptom—a symptom of irreconcilable constraints and pressures deeper inside a system; a pointer to systemic trouble further upstream.

The Bad Apple Theory II: Error classification systems

In order to lead people (e.g. investigators) to the sources of human error as inspired by Fitts and Jones '47, a number of error classification systems have been developed in aviation (e.g. the Threat and Error Management Model (e.g. Helmreich *et al.*, 1999; Helmreich, 2000) and the Human

Factors Analysis and Classification System (HFACS, Shappell & Wiegmann, 2001)). The biggest trap in both error methods is the illusion that classification is the same as analysis. While classification systems intend to provide investigators more insight into the background of human error, they actually risk trotting down a garden path toward judgments of people instead of explanations of their performance; toward shifting blame higher and further into or even out of organizational echelons, but always onto others. Several false ideas about human error pervade these classification systems, all of which put them onto the road to The Bad Apple Theory.

First, error classification systems assume that we can meaningfully count and tabulate human errors. Human error "in the wild", however—as it occurs in natural complex settings—resists tabulation because of the complex interactions, the long and twisted pathways to breakdown and the context-dependency and diversity of human intention and action. Labeling certain assessments or actions in the swirl of human and social and technical activity as causal, or as "errors" and counting them in some database, is entirely arbitrary and ultimately meaningless. Also, we can never agree on what we mean by error:

- Do we count errors as *causes* of failure? For example: This event was due to human error.
- Or as the *failure itself*? For example: The pilot's selection of that mode was an error.
- Or as a *process*, or, more specifically, as a departure from some kind of standard? This may be operating procedures, or simply good airmanship. Depending on what you use as standard, you will come to different conclusions about what is an error.

Counting and coarsely classifying surface variabilities is protoscientific at best. Counting does not make science, or even useful practice, since interventions on the basis of surface variability will merely peck away at the margins of an issue. A focus on superficial similarities blocks our ability to see deeper relationships and subtleties. It disconnects performance fragments from the context that brought them forth, from the context that accompanied them; that gave them meaning; and that holds the keys to their explanation. Instead it renders performance fragments denuded: as uncloaked, context-less, meaningless shrapnel scattered across broad classifications in the wake of an observer's arbitrary judgment.

Second, while the original Fitts and Jones legacy lives on very strongly in human factors (for example in Norman (1994) who calls technology something that can make us either smart or dumb), human error classification systems often pay little attention to systematic and detailed nature of the connection between error and people's tools. According to Helmreich (2000), "errors result from physiological and psychological limitations of humans. Causes of error include fatigue, workload, and fear, as well as cognitive overload, poor interpersonal communications, imperfect information processing, and flawed decision making" (p. 781). Gone are the systematic connections between people's assessments and actions on the one hand, and their tools and tasks on the other. In their place are purely human causes—sources of trouble that are endogenous; internal to the human component. Shappell and Wiegmann, following the original Reason (1990) division between latent failures and active failures, merely list an undifferentiated "poor design" only under potential organizational influences—the fourth level up in the causal stream that forms HFACS. Again, little effort is made to probe the systematic connections between human error and the engineered environment that people do their work in. The gaps that this leaves in our understanding of the sources of failure are daunting.

Third, Fitts and Jones remind us that it is counterproductive to say what people failed to do or should have done, since none of that explains why people did what they did (Dekker, 2001). With the intention of explaining why people did what they did, error classification systems help investigators label errors as "poor decisions", "failures to adhere to brief", "failures to prioritize attention", "improper procedure", and so forth (Shappell & Wiegmann, 2001, p. 63). These are not explanations, they are judgments. Similarly, they rely on fashionable labels that do little more than saying "human error" over and over again, re-inventing it under a more modern guise:

- Loss of CRM (Crew Resource Management) is one name for human error—the failure to invest in common ground, to share data that, in hindsight, turned out to have been significant.
- Complacency is also a name for human error—the failure to recognize the gravity of a

situation or to adhere to standards of care or good practice.

- Non-compliance is a name for human error—the failure to follow rules or procedures that would keep the job safe.
- Loss of situation awareness is another name for human error—the failure to notice things that in hindsight turned out to be critical.

Instead of explanations of performance, these labels are judgments. For example, we judge people for not noticing what we now know to have been important data in their situation, calling it their error—their loss of situation awareness.

Fourth, error classification systems typically try to lead investigators further up the causal pathway, in search of more distal contributors to the failure that occurred. The intention is consistent with the organizational extension of the Fitts and Jones '47 premise (see Maurino *et al.*, 1995) but classification systems quickly turn it into re-runs of The Bad Apple Theory. For example, Shappell & Wiegmann (2001) explain that "it is not uncommon for accident investigators to interview the pilot's friends, colleagues, and supervisors after a fatal crash only to find out that they 'knew it would happen to him some day'." (p. 73) HFACS suggests that if supervisors do not catch these ill components before they kill themselves, then the supervisors are to blame as well. In these kinds of judgments the hindsight bias reigns supreme (see also Kern, 1999). Many sources show how we construct plausible, linear stories of how failure came about once we know the outcome (e.g. Starbuck & Milliken, 1988), which includes making the participants look bad enough to fit the bad outcome they were involved in (Reason, 1997). Such reactions to failure make after-the-fact data mining of personal shortcomings—real or imagined—not just counterproductive (sponsoring The Bad Apple Theory) but actually untrustworthy. Fitts' and Jones' legacy says that we must try to see how people—supervisors and others—interpreted the world from their position on the inside; why it made sense for them to continue certain practices given their knowledge, focus of attention and competing goals. The error classification systems do nothing to elucidate any of this, instead stopping when they have found the next responsible human up the causal pathway. "Human error", by any other label and by any other human, continues to be the conclusion of an investigation, not the starting point. This is the old view of human error, re-inventing human error under the guise of supervisory shortcomings and organizational deficiencies. HFACS contains such lists of "unsafe supervision" that can putatively account for problems that occur at the sharp end of practice. For example, unsafe supervision includes "failure to provide guidance, failure to provide oversight, failure to provide training, failure to provide correct data, inadequate opportunity for crew rest" and so forth (Shappell & Wiegmann, 2001, p. 73). This is nothing more than a parade of judgments: judgments of what supervisors failed to do, not explanations of why they did what they did, or why that perhaps made sense given the resources and constraints that governed their work. Instead of explaining a human error problem, HFACS simply re-locates it, shoving it higher up, and with it the blame and judgments for failure. Substituting supervisory failure or organizational failure for operator failure is meaningless and explains nothing. It sustains the fundamental attribution error, merely directing its misconstrued notion elsewhere, away from front-line operators.

In conclusion, classification of errors is not analysis of errors. Categorization of errors cannot double as understanding of errors. Error classification systems may in fact reinforce and entrench the misconceptions, biases and errors that we always risk making in our dealings with failure, while giving us the illusion we have actually embraced the new view to human error. The step from classifying errors to pursuing culprits appears a small one, and as counterproductive as ever. In aviation, we have seen The Bad Apple Theory at work and now we see it being re-treaded around the wheels of supposed progress on safety. Yet we have seen the procedural straight-jacketing, technology-touting, culprit-extraditing, train-and-blame approach be applied, and invariably stumble and fall. We should not need to see this again. For what we have found is that it is a dead end. There is no progress on safety in the old view of human error.

People create safety

We can make progress on safety once we acknowledge that people themselves create it, and we begin to understand how. Safety is not inherently built into systems or introduced via isolated technical or procedural fixes. Safety is something that people create, at all levels of an operational organization (e.g. AMA, 1998; Sanne, 1999). Safety (and failure) is the emergent property of entire systems of people and technologies who invest in their awareness of potential pathways to breakdown and devise strategies that help forestall failure. The decision of an entire airline to no longer accept NDB approaches (Non-Directional Beacon approaches to a runway, in which the aircraft has no vertical guidance and rather imprecise lateral guidance) (Collins, 2001) is one example of such a strategy; the reluctance of airlines and/or pilots to agree on LASHO—Land And Hold Short Operations—which put them at risk of traveling across an intersecting runway that is in use, is another. In both cases, goal conflicts are evident (production pressures versus protection against known or possible pathways to failure). In both, the trade-off is in favor of safety. In resource-constrained systems, however, safety does not always prevail. RVSM (Reduced Vertical Separation Minima) for example, which will make aircraft fly closer together vertically, will be introduced and adhered to, mostly on the back of promises from isolated technical fixes that would make aircraft altitude holding and reporting more reliable. But at a systems level RVSM tightens coupling and reduces slack, contributing to the risk of interactive trouble, rapid deterioration and difficult recovery (Perrow, 1984). Another way to create safety that is gaining a foothold in the aviation industry is the automation policy, first advocated by Wiener (e.g. 1989) but still not adopted by many airlines. Automation policies are meant to reduce the risk of coordination breakdowns across highly automated flight decks, their aim being to match the level of automation (high, e.g. VNAV (Vertical Navigation, done by the Flight Management System); medium, e.g. heading select; or low, e.g. manual flight with flight director) with human roles (pilot flying versus pilot not-flying) and cockpit system display formats (e.g. map versus raw data) (e.g. Goteman, 1999). This is meant to maximize redundancy and opportunities for double-checking, capitalizing on the strengths of available flightdeck resources, both human and machine.

When failure succeeds

People are not perfect creators of safety. There are patterns, or mechanisms, by which their creation of safety can break down—mechanisms, in other words, by which failure succeeds. Take the case of a DC-9 that got caught in windshear while trying to go around from an approach to Charlotte, NC, in 1994 (NTSB, 1995). Charlotte is a case where people are in a double bind: first, things are too ambiguous for effective feedforward. Not much later things are changing too quickly for effective feedback. While approaching the airport, the situation is too unpredictable, the data too ambiguous, for effective feedforward. In other words, there is insufficient evidence for breaking off the approach (as feedforward to deal with the perceived threat). However, once inside the situation, things change too rapidly for effective feedback. The microburst creates changes in winds and airspeeds that are difficult to manage, especially for a crew whose training never covered a windshear encounter on approach or in such otherwise smooth conditions.

Charlotte is not the only pattern by which the creation of safety breaks down; it is not the only mechanism by which failure succeeds. For progress on safety we should de-emphasize the construction of cause—in error classification methods or any other investigation of failure. Once we acknowledge the complexity of failure, and once we acknowledge that safety and failure are emerging properties of systems that try to succeed, the selection of causes—either for failure or for success—becomes highly limited, selective, exclusive and pointless. Instead of constructing causes, we should try to document and learn from patterns of failure. What are the mechanisms by which failure succeeds? Can we already sketch some? What patterns of breakdown in people's creation of safety do we already know about? Charlotte—too ambiguous for feedforward, too dynamic for effective feedback—is one mechanism by which people's investments in safety are outwitted by a rapidly changing world. Understanding the mechanism means becoming able to retard it or block it, by reducing the mechanism's inherent coupling; by disambiguating the data that fuels its progression from the inside. The contours of many other patterns, or mechanisms of failure, are beginning to stand out from thick descriptions of accidents in aerospace, including the normalization of deviance (Vaughan, 1996), the going sour progression (Sarter & Woods, 1997), practical drift (Snook, 2000) and plan continuation (Orasanu *et al.*, in press). Investing further in these and other insights will represent progress on safety. There is no efficient, quick road to understanding human error, as error classification methods make us believe. Their destination will be an illusion, a retread of the old view. Similarly, there is no quick safety fix, as the punishment of culprits would make us believe, for systems that pursue multiple competing goals in a resource-constrained, uncertain world. There is, however, percentage in opening the black box of human performance—understanding how people make the systems they operate so successful, and capturing the patterns by which their successes are defeated.

Acknowledgements

The work for this paper was supported by a grant from the Swedish Flight Safety Directorate and its Director Mr. Arne Axelsson.

References

American Medical Association (1998). *A tale of two stories: Contrasting views of patient safety*. Report from a workshop on assembling the scientific basis for progress on patient safety.

Chicago, IL: National Patient Safety Foundation at the AMA.

Collins, R. L. (2001). No NDB. *Flying, 128*(1), 18.

Dekker, S. W. A. (2001). The disembodiment of data in the analysis of human factors accidents. *Human Factors and Aerospace Safety, 1*(1), 39-57.

Fitts, P. M., & Jones, R. E. (1947). Analysis of factors contributing to 460 'pilot error' experiences in operating aircraft controls. *Memorandum Report TSEAA-694-12,* Aero Medical Laboratory, Air Material Command, Wright-Patterson Air Force Base, Dayton, Ohio, July 1, 1947.

Goteman, O. (1999). Automation policy or philosophy? Management of automation in the operational reality. In S. W. A. Dekker & E. Hollnagel (Eds.), *Coping with computers in the cockpit,* pp. 215-222. Aldershot, UK: Ashgate Publishing Co.

Hawkins, F. (1987). *Human factors in flight.* Aldershot, UK: Avebury.

Helmreich, R. L., Klinect, J. R., & Wilhelm, J. A. (1999). Models of threat, error and response in flight operations. In R. S. Jensen (Ed.), *Proceedings of the tenth international symposium on aviation psychology*. Columbus, OH: The Ohio State University.

Helmreich, R. L. (2000). On error management: Lessons from aviation. *BMJ,* 320, 745-753.

Kern, T. (1999). *Darker shades of blue: The rogue pilot*. New York, NY: McGraw-Hill.

Maurino, D. E., Reaon, J. T., Johnston, N., & Lee, R. B. (1995). *Beyond aviation human factors*. Aldershot, UK: Ashgate Publishing Co.

National Transportation Safety Board (1995). *Flight into terrain during missed approach. USAir Flight 1016, DC-9-31, Charlotte/Douglas International Airport, Charlotte, NC, 7/2/94* (NTSB Rep. No. AAR/95/03). Washington, DC: NTSB.

Norman, D. A. (1994). Things that make us smart: Defending human attributes in the age of the machine. New York: Perseus Press.

Orasanu, J., Martin & Davison (in press). Sources of decision error in aviation. In G. Klein and E. Salas (Eds.), *Applications of naturalistic decision making*. Mahwah, NJ: Lawrence

Erlbaum Associates.

Perrow, C. (1984). *Normal accidents: Living with high-risk technologies*. New York, NY: Basic books.

Reason, J. T. (1990). *Human error*. Cambridge, UK: Cambridge University Press.

Reason, J. T. (1997). *Managing the risks of organizational accidents.* Aldershot, UK: Ashgate Publishing Co.

Reason, J. T. (2000). Grace under fire: Compensating for adverse events in cardiothoracic surgery. *Paper presented at the 5th conference on naturalistic decision making,* Tammsvik, Sweden. May, 2000.

Sanne, J. M. (1999). *Creating safety in air traffic control*. Lund, Sweden: Arkiv.

Sarter, N. B., & Woods, D. D. (1997). Teamplay with a powerful and independent agent: Operational experiences and automation surprises on the Airbus A-320. *Human Factors, 39*(4), 553-569.

Shappell, S. A., & Wiegmann, D. A. (2001). Applying Reason: the human factors analysis and classification system (HFACS). *Human Factors and Aerospace Safety, 1*(1), 59-86.

Simon, H. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.

Snook, S. A. (2000). *Friendly fire: The accidental shootdown of US Blackhawks over Northern Iraq*. Princeton, NJ: Princeton University Press.

Starbuck, W. H., & Milliken, F. J. (1988). Challenger: Fine-tuning the odds until something breaks. *Journal of Management Studies, 25*(4), 319-340.

Vaughan, D. (1996). *The Challenger lauch decision: Risky technology, culture and deviance at NASA*. Chicago, IL: University of Chicago Press.

Wiener, E. L. (1989). *Human factors of advanced technology ("glass cockpit") transport aircraft* (NASA contractor report No. 177528). Moffett Field, CA: NASA Ames Research Center.

Woods, D. D., Johannesen, L. J., Cook, R. I., & Sarter, N. B. (1994). *Behind human error: Cognitive systems, computers and hindsight.* Dayton, OH: CSERIAC.

